Hermeneutics and Information Representation

Matthew Chalmers

Computing Science, Glasgow University, United Kingdom

Abstract

By drawing from semiology, epistemology and philosophical hermeneutics, we discuss the way CSCW models data, situation and activity—information representation. We point out similarities between discourse in hermeneutics and in the anthropology and sociology that predominates in CSCW, and propose that a hermeneutic perspective offers a unifying view on the social science and computer science within CSCW. We discuss formalisation, adaptation, and objectivity in our theories, methodologies and implementations, and offer collaborative filtering and its extension, the path model, as examples of practical approaches to information representation that show, support and adapt with activity in a hermeneutic style.

Introduction

The key issue in CSCW, according to Schmidt and Bannon (1992), is "the problem of how to support the ongoing dynamic articulation of distributed activities and the cooperative management of the mechanisms of interaction themselves." They focus on articulation work, and the resolution of inconsistencies between the real world and the computer's model. As was later quoted in (Bannon and Bødker, 1997)

> Cooperative work is not facilitated simply by the provision of a shared database, but requires the active construction by the participants of a common information space where the meanings of the shared objects are debated and resolved, at least locally and temporarily. Objects must

thus be interpreted and assigned meaning, meanings that are achieved by specific actors on specific occasions of use.

Bannon and Bødker pointed out CSCW's lack of focus on the common information space (CIS) and how, in the intervening five years, "the work involved in both putting information in common, and in interpreting it, has not been sufficiently recognised." This work is a primary topic of this paper. We see it as being highly co–dependent on the databases and programs used in this sharing and interpretation. The design of these information representations, and their status as symbolic or semiological artifacts, are central to CSCW and HCI. As in (Dourish 1997), we consider that "HCI's concern with the role of systems in supporting people and work must, then, be based on some understanding of the representations through which computer systems operate." Dourish focuses on the ontological status of these representations and the way that an information representation reflects an intensional stance on the world.

These are issues of concern to epistemology and semiology. We see these as fields of discourse relatively underutilised in CSCW, and throughout this paper we apply them in discussing CSCW theory and practice. We take a pragmatic epistemological stance where there is no meaning for a symbol independent of human interpretation. We make particular reference to the philosophical hermeneutics of Gadamer (1989) to which (Warnke, 1987) offers an introduction. To those wishing a more general text, we suggest (Coyne, 1995) and (Grondin, 1994).

In Gadamer's hermeneutics, the meaning of a representation or symbol is open to interpretation. It is not an absolute, defined independently of other things and symbols. This is in contrast to positivism, which sees a symbol as logically and historically naming a thing in the real world. Positivism underlies traditional informatics' focus on fixed representations of recorded phenomena, independent of later use, so that computer 'memory' tends to exclude representations of individuals' context and situation. Dialogue and the "dialectic nature of the CIS," as Bannon and Bødker put it, are not part of mainstream informatics' theories and designs.

Taking a wider view, we see this paper as following in the footsteps of those within hermeneutics that seek a view of knowledge and inquiry that unifies the perspectives of the social and natural sciences. One might see such an intellectual tool as having two sides (or perhaps as being double–edged). One extensional side would relate to our discourse on CSCW in general i.e. our theories and general approaches. The other intensional side would relate to particular methodologies and implementations. Informatics needs such an implement to cut free from a reliance on positivism in its theories and systems.

Speaking, Writing and Distanciation

Schmidt and Bannon, in their discussion of putting information in common, suggested that we can move beyond the simple database to the CIS where "the producer and consumer [of information] consciously make an effort to understand each other's context." In effect, they are moving towards dialogical communication, where breakdowns in meaning and understanding are debated and resolved. Each speaker's dialogical actions are made to an explicit audience, can be rich in ostensive reference, and take place in a shared context of use. Speakers can use indexical forms of language, which are conditional or reliant on situation for significance. Suchman contrasts indexical expressions with, for example, "definite noun phrases whose meaning is claimed to be specifiable in objective, or context–independent terms" (1987:59).

Perhaps the best known and most used databases are those that underlie the information retrieval systems on the World Wide Web such as the AltaVista search engine. The database volume and the number of users of such systems makes direct dialogue impractical. Producer and consumer are in a situation far more akin to writing than speech, and so would be better described as 'author' and 'reader.' According to hermeneuticists such as Paul Ricoeur (1981), writing shows four characteristic traits of *distanciation*. The author's meaning is inscribed in the text, a fixed, finite and external representation, the text is dissociated from the mental intention of the author, can only display non–ostensive references, and, instead of a known audience directly apparent to a speaker, the range of potential readers is unlimited. These four traits taken together constitute the text's 'objectivity' (p. 210) or semantic autonomy (p. 37).

The distanciation of the text is thus related to its potential use in communication. Continuing the quotation of Suchman, and retaining her emphasis:

> But the *communicative* significance of a linguistic expression is always dependent upon the circumstances of its use. A formal statement not of what the language means in relation to any context, but of what the language–user means in relation to some particular context, requires a description of the context or situation of the utterance itself. And every utterance's situation comprises an indefinite range of possibly relevant features.

Hermeneutic theory is based on accepting the effect of this indefinite, inevitable and infinitely detailed situational background. Any use of a symbol in a human activity carries with it assumptions, abbreviations, applicability, the people involved, the other information that they share as part of their current activity, their organisational structure and practices, and so on, endlessly. Each language–user's interpretation involves their understanding of the original utterance's content and context, as well as understanding of other contexts of use of the utterance and contained words and symbols i.e. their position in the reader's language or tradition. In hermeneutic terms, this involves the individual's *pre–understanding* or *prejudice*. These terms emphasise an understanding effective

prior to, or without, conscious understanding. Prejudice here also means 'pre–judgement' i.e. a preliminary judgement that is always open to refutation or adaptation. To reject prejudice as always false or misleading is, in Gadamer's view, simply a "prejudice against prejudice" (1989:240).

Formalisation and Action

Schmidt (1997) recently discussed the tension between the detailed contingency of situated action and the process of generalisation over the specificities of situation necessary to create formal constructs. He points out their crucial coordinative and communicative significance, and the way that formal constructs "would be of only marginal utility if they were not inscribed upon artifacts." Such an artifact is for him a shared space with a role:

> [To] give permanence to the protocol for which it stands proxy in the sense that it conveys the stipulations of the protocol in a situation–independent manner. […] Written artifacts can at any time be mobilized as a referential for clarifying ambiguities and settling disputes: 'while interpretations vary, the word itself remains as it always was.' They are, for all practical purposes, unceasingly publicly accessible.

We see formal constructs as texts whose design necessarily involves distanciation. Schmidt puts forward a strong case as to the variability of their normative effect, based on their underspecification of situation. This effect may range from script–like sequentiality of actions to flexibly interpreted, map–like resources for situated action.

Although the inscription remains 'as it always was,' its interpretation by subsequent readers may vary. The intentions of the original author are not privileged. Each interpretation is itself situated in a context of activity, person and history. According to Gadamer, interpretation is most usefully seen as an interaction between the activity, context and prejudice (the 'horizon') of the reader, and the content, context and background of the information i.e. the horizon provided by the text. Gadamer holds that every reading or hearing of a text constitutes an act of giving meaning to it through a 'fusion of horizons.' Interpretation is based on prejudice, which includes assumptions implicit in the language that the person uses. That language in turn is learned through experiences of interpretation. The individual and their prejudice are changed through the use of language, and the language changes through its use by individuals. A new word or experience is understood in relation to, and within, language and history. This endless process of seeing the part in and through the whole is the *hermeneutic circle*.

Interpretation of a formal construct may thus potentially give rise either to ambiguities or to clarifications, and may create disputes or settle them. Formal constructs, as with all texts or text analogues, show sometimes deliberate distanciation and decreased emphasis on situational detail. These attributes bring

openness to interpretation and varied semiological use—Schmidt's 'underspecification.'

Another important aspect of hermeneutic theory here is its scope. Ricoeur uses *meaningful action* (or 'meaning–bearing' action) as the definition of the object of the social sciences. Here he states that he borrows from Weber and, by making the analogy between text and action, "enables the theory of interpretation to be freed from the constraints of discourse and writing, and extended to the whole field of the social sciences" (p. 37). Application of hermeneutic theory of interpretation is not limited to speech and writing *per se*, but instead "the intermediary link between the model of the text and social phenomena is constituted by the notion of semiological systems" (p. 218). One can choose to use anything and any combination of media to communicate. It is this interpretative choice or reaction that creates significance, and so any action in any medium can be taken as meaningful or significant, and hence as a symbol.

If we treat meaningful action as being semiological then the structuralist models (Saussure, 1983) that hermeneutics builds on afford a kind of explanation of action far from strict causality or normative sequentiality: *correlative* description. Meaning is based in recurrent—statistically consistent—patterns of symbol use. Or, as in (Wittgenstein, 1958), "the meaning of a word is its use in the language."

Statistically consistent patterns of symbol use, that span media and that adapt for each individual with interaction by means of the hermeneutic circle, are also at the core of evolutionary linguistics (Hurford, 1998). This field, sprouting vigorously in recent years, examines how language begins and evolves. It grounds language in social cognition, and suggests that evolutionary processes can lead to the formation of complex structures such as a consistent grammar and syntax within a community, without recourse to a shared and 'prewired' Chomskyan protolanguage in individuals' brains. It holds that these same processes of interaction, adaptation and selection create the phonology i.e. the patterns in phenomena such as sound and shape that are used as symbols within a language. These processes also determine which phenomena are used and combined in symbols and language. To quote from the book's introduction:

> Implicit in this work and in the work of much historical linguistics that still goes forward is a view of language as a property of a group rather than of an individual. […] For Saussure, 'language is not complete in any speaker; it exists only within a collectivity…only by virtue of a sort of contract signed by members of a community.'

Each individual, through personal experience of the 'give and take' of interacting with others, and through the hermeneutic circle, builds their own ongoing knowledge of semiological practice: Saussure's *langue*. The full richness of personal experience is woven into the knowledge of the use of language. When we, through scientific observational procedure more than personal experience, inscribe and formalise such statistical generalisations, we do not represent all of

that richness. Such a written formal construct only approximates the *langue* of the observed speakers at the time of observation and, as time goes on, distanciation increases between their dynamically evolving language and our static representation of it. In other words, the hermeneutic circle is broken and 'ontological drift' (Robinson and Bannon, 1991) occurs.

Formal constructs are often seen as 'meta–objects' that stand apart from the language they describe, but this contradicts the way the hermeneutic circle involves the continual integration of the part with the whole. Structuralism's *configurational* view, where a symbol's meaning depends on its differences and similarities with the rest of the configuration of symbols, and a change in use and meaning of any object may affect every other, is also relevant here. When one tries to formalise aspects of the world and form a metalanguage, one is still making and using symbols, and one finds oneself slipping back down to language again i.e. as soon as people start to use a metalanguage, it becomes language. They use it in their everyday work or activity, and that activity is carried out in natural language. Meta–symbolic information has no position exterior to the system of language.

Objectivity

The language understood by a person or a community is large but finite, and has a remarkable ability to adapt with use. It is an open system. Our formal constructs are also finite, but generally they are fixed and closed. In designing a formalisation, a set of symbols is chosen: which objects (including meta–objects) are included and which phenomena are involved in their description. As we have seen, the separate treatment of meta–information is one way in which a more easily managed set of symbols is obtained. Another way is to restrict the phenomena involved in a representation, for example representing a text by the words it contains but ignoring its context of use.

Traditional scientific observational procedure attempts to create manageable, manipulable representations of the world. It aims to reach objectivity, semantic autonomy and independence from individual interpretation. In the paradigm of the natural sciences, a situation can be described formally and precisely in terms of scientific laws, and subsequent behaviour will follow causal patterns with unerring consistency. Actions and behaviours are experimentally repeatable by anyone, anywhere, and in any situation. In this case, we see no use in attempting to represent and communicate the infinitely detailed context of person and situation that is characteristic of the social sciences. Our 'texts' can be formal and finite, yet useful: they are 'data.' They afford repeatability of action, which brings the predictive power of the natural sciences. This is in contrast to the social sciences that find concise, static formalisms less useful because of situational dependence. Here, repeatability and predictability are relatively weak. Our texts

are not data, but 'information.' The language or discourse of the natural sciences seems like writing, while that of the social sciences is more like speech.

We are increasingly aware of the limits of objectivity, understanding scientific theories as impermanent and pragmatic constructs that may be falsified and overturned. They are interpretive frameworks dependent on observational and practical goals. Wittgenstein (1958) showed that proof is a 'language game' whose truth, as with all our natural language, is determined by our own social use rather than axiomatic deduction. The foundations of the natural sciences are semiological, social, and subjective. Wegner (1997) recently showed (via Turing) that the unbounded possibilities for human action make interactive systems necessarily incomplete i.e. not reducible to behaviour describable by formal algorithms. Wegner's proof is a confirmation of the criticisms of Wittgenstein (and Gödel), expressed within the formal language of informatics. Informatics shows itself as semiological.

Writers such as Rorty (1979) have put forward a related view, that hermeneutics overcomes unnecessary distinctions between the knowledge of the natural and social sciences, and that we should look towards a pragmatic basis for all inquiry. Following Ricoeur, we concentrate instead on correlative explanations. Induction from patterns of correlation can be made, as is the tradition of science, but we should treat this induction as preliminary judgement or prejudice that, echoing Gadamer, may later be refuted or adapted. 'Objectivity' comes from distanciation: representation is fixed, dissociated from intention, and only displays universally shared references. In the case of the natural sciences we emphasise distanciation, using maximal correlative consistency (in place of absolute causality) and minimal dependence on situation. In the social sciences, we de–emphasise distanciation, and so correlative consistency may decrease and situational dependence increase. There is no dichotomy between the two, as objectivity is not absolute. Instead we see degrees of distanciation.

Similarly, we see no dichotomy between data and information. Abstraction, formalisation and inscription can be used as a means to move from information to data, and to increase predictability. What is then predicted, however, is that the same generalisations will recur—the same correlations, averages and distributions. We describe or explain the 'average situation,' but no actual situation. Our data may fit with our model of the 'average person,' but our systems suit no–one in particular.

Even if this unity of the sciences and universality of hermeneutics is accepted, an important weakness remains. If the written explanation or account is fixed, and taken out of the hermeneutic circle, then the ontological drift characteristic of writing can begin.

We now begin to summarise and tease apart some of the issues and phenomena we have discussed so far. We will use the following list of landmarks or reference points to guide our discussion, in later sections, of schemes of interpretation and

representation in CSCW:

- *Phenomena:* Which phenomena are represented or emphasised? Examples are the people involved, the context of work activity, the content of objects such as documents, and meta–information. How are phenomena interpreted: formally and objectively, or flexibly and subjectively?
- *Sharing:* Is the representation a shared resource that the entire community can draw from—whether as map or script?
- *Interaction:* Does each moment of use involve a summary or overview of the entire representation? Or do we consider only the most relevant details or subparts of the representation? If the latter, does this require explicit declaration and formalisation of relevance, or is this more passively determined e.g. by extrapolating from past activity?
- *Adaptation:* Is the representation fixed *a priori* or does the system adapt with every use? How are objects added to or deleted from the representation?
- *Configurationality:* Does the representation of each object depend on its differences and similarities with the rest of the configuration of symbols?

Before continuing, we briefly reflect on this list. No doubt incomplete, the list is offered as a rough map to initiate discussion and sharing of ideas, rather than as a script or formal checklist. While summarising and extending earlier discussion, it is preliminary and can be expected to be adapted in the light of future discussion and experience. Also, it should not be taken as a strict partition of design concerns; clearly the listed points overlap with and depend on each other.

This framework has previously been used to discuss and compare workflow, information retrieval, collaborative filtering and the author's extension of collaborative filtering, the path model (Chalmers, 1999a). We will discuss the latter two here, as examples serving our argument that an approach to theory and design is beginning to take shape, that is complementary to mainstream CSCW and with a standpoint nearer to hermeneutics. First, however, we use the framework in discussing two related areas of CSCW, from the observational side and from the system side, which may not encompass the mainstream but perhaps exemplify its current state. We review ethnomethodology and open implementation, especially with regard to the issue of distanciation.

Ethnomethodology and Open Implementation

We have so far pointed out a number of parallels between discourse in CSCW based on sociology and anthropology, and discourse in hermeneutics. As was noted in (Giddens, 1995), Garfinkel's ethnomethodology has strong links with hermeneutics despite the former's focus on the more intensional task of generating an empirical research programme and the latter's expression in abstract philosophy. Giddens also points out that the two fields share historical roots in the Wittgenstein's language games. The central phenomenon of

ethnomethodology is the performance or representation of work that reveals work's organisation. It focuses on meta–level phenomena: the rules and patterns of social conduct that are an essential and emergent part of everyday activity, and not external or distant from it. Meta–language is part of language.

Giddens criticises Garfinkel, or those immediately influenced by him, for not addressing their "unsophisticated" epistemological stance. He raises, for instance, the imbalance between ethnomethodology's concentration on the phenomenon of agency in the production of society, and not on the complementary, structural analysis of society's reproduction e.g. the act of speaking grammatically as reproducing the rules used to generate the utterance. Another issue is the emphasis on public accountability when describing conversations, while the goals and motives of individuals are less often used.

While ethnomethodologists work and analyse a setting, these rules and patterns are a shared resource that the entire community has at hand. During this time they can be discussed and adapted, but it is impractical for analysis to continue indefinitely. Then, drift and distanciation can begin.

This bounded adaptability is most obvious when an account of a setting is used in the design of new technologies for that setting. The tradition of ethnomethodology is in analysing practice, rather than "inventing the future" (Button and Dourish, 1996). In bridging from ethnomethodological studies of use and critique of system design, over to the practice of system design, we see what Button and Dourish call the 'paradox of technomethodology.' The focus on detailed context, on the individuals involved and on the moment–by–moment organisation of action, clashes with the design of new technologies that will transform patterns of action, and invalidate the ethnomethodological account. While the shared world described by ethnomethodology can influence the designer's private work, this description is of a past world—from a time before the program was introduced.

Button and Dourish suggest that the best way to learn from ethnomethodology is to change system design practice at a fundamental level, and offer as an example the use of computational reflection and open implementations (Dourish, 1995). Open implementations are so called because they reveal policies of deep system behaviour. In a related paper, Bentley and Dourish (1995) put forward an incremental customisation approach, based on open implementations and treating a CSCW system as "a *medium* through which collaborative work occurs, rather than an embodiment of *mechanisms* representing perceived regularities in collaborative activity". They emphasise the power that flexible interpretation of activity brings, and see customisation as "not simply a method for individuals to adapt technology to meet their own needs; it is, fundamentally, a means by which users can *construct* their working patterns." Users can then change how they work by design.

Bentley and Dourish see a 'customisation gulf' that lies between currently

available superficial customisation and user adaptation of deep system behaviour. The gulf involves two characteristic problems: the level of customisation*,* which is generally shallow, and the language of customisation, knowledge of which is generally outwith the skills users have or are willing to obtain. Although the term refers more to explicitly controlled system adaptation, we see the customisation gulf as being roughly synonymous with the hermeneutic circle. When it is bridged, people adapt system behaviour in accordance with their use.

Bentley and Dourish also put forward computational reflection as a means to construct open implementations. Here, the emphasis is on having a part of the system causally connected to the other parts, and that part operates at a 'meta–level.' The meta–level provides a continually available account or description of the system's behaviour, similar to ethnomethodological accounts, and it adapts system functionality in accordance with ongoing activity. In theory, reflective systems have unbounded adaptability, via self–generated and self–modifying code, but in practice the range of policies is generally small and fixed. This is partly because of the complexity of designing meta–object protocols and self–modifying code, and partly to ensure that the account of system behaviour is manageable for users.

Designing a meta–object protocol is a very difficult programming task. Part of the complexity of reflective programs is their internal connectedness and configurationality. Also, the designer must predict what, at all levels of abstraction, will be required by users in the future. When the designer finishes writing and the program is in use then—just as distanciation starts when ethnomethodological observation stops—these predictions do not change or adapt. When the aspects of the world represented are significantly varied and unpredictable, such as in people's work, then the limits of bounded choice and adaptation may constrain that work. Open implementation cannot reveal the entire architecture, but accounts can be used to simplify and tailor the representation of the architecture and system policies to users. Users' interpretation of the system then involves this interface, but this relies on the designer's interpretation of the users.

The distanciation of reflective programs limits the bridging of the customisation gulf. This limit would disappear if the theoretically unbounded adaptability of reflective systems is achieved in practice. At present, however, the paradox of technomethodology is due to the hermeneutic circle being broken—in the same way—in the practices of both computational reflection and ethnomethodology.

Data and Information, Space and Place

We can now see part of the reason for the elusiveness of the perhaps idealised common information space, discussed at the beginning of this paper. The CIS

requires a degree of dialogical communication not afforded by tools such as technomethodology. It is their inherent distanciation that causes this.

If we take a few steps on, via a paper related to those cited in the previous section, we can obtain another useful perspective on these problems. In (Harrison and Dourish, 1996), the notions of *space* and *place* were distinguished. The former seems closer to what we often have now in our attempts to build a CIS, while the latter is closer to the ideal. A *space* here need not be a physical space, of course, but also refers to virtual environments and media spaces i.e. information spaces. A *place* involves more than a three dimensional structure or a configuration of information objects. It also involves what they call 'appropriate behavioural framing,' the emergent patterns of human behaviour and interaction that offer understandings of the space. It might be, therefore, that Harrison and Dourish would prefer 'common information place' as the name for our ideal.

They focus on information spaces that employ aspects of physical spaces in order to support interaction. They suggest that naïve mimicry of the physical world does not significantly help in making a useful place out of a sterile information space. One critical factor, they say, is "support for adaptation and appropriation of the technology by user communities." This leads to the development of a "communally-held sense of appropriate behaviour, and a context for engaging in and interpreting action." Again we see the issue of the distance between the information representation's designers and its users. Even the technologically advanced methods of open implementation are ultimately strongly inhibited in affording users' adaptation and appropriation.

Harrison and Dourish interpret issues of technology design and information representation in terms of space and place, and build up a complex set of terms and concepts, such as 'spaceless places,' in discussing these issues. Our approach pursue relies on what we consider as a simpler but wider view. We follow the example of Ricoeur mentioned earlier, treating space as a medium for significant action, adding it to the media that collectively form the semiological system i.e. we consider spatial structure and action as being part of language.

We consider 'space' to be one of the physical phenomena which has the potential to be used as a symbol, and so consider 'place' to be the symbol in language, given meaning by its statistical patterns of recurrence in human use. The meaning of a space is its use in the language. Harrison and Dourish's principle "Space is the opportunity; place is the understood reality" now can be seen as a relatively straightforward restatement of hermeneutic concepts. Their 'appropriate behavioural framing' is 'grammatical language,' semiological activity conforming to social norms. We can similarly handle the distinction between data and information alluded to earlier with regard to objectivity. We see 'data' as an informational phenomenon that has the potential to be used symbolically, and 'information' as the symbol in language.

This treatment of symbols, spanning phenomena or media such as space, data

and text, also makes the 'complex forms' of Harrison and Dourish conceptually simpler. The first are 'spaceless places,' such as Usenet news groups, that are navigated and used by means of relationships that are non-spatial but that nevertheless support "the tension between connectedness and distinction which leads to placefulness." As they put it, different social norms make for different places without an underlying notion of space. We see them as different languages or semiological systems. Whether or not spatial notions underlie them is unproblematic.

Their second class of complex forms, hybrid spaces, are also easier to conceptualise from this viewpoint. In this case, physical and virtual spaces merge or overlap. Actions in physical space may be 'projected' into a media space by means of cameras. While it may be difficult for the people involved to handle symbols that have been projected into a new context, conceptually it appears simple to consider these actions as sharing a common set of symbols, involving these various media. In fact, we find it difficult to consider spaces that are *not* hybrid. Actions, words, gestures, concepts and goals that are not particular to one workplace or information space are inevitably involved, and span such spaces because the same people are involved in them all. Media spaces, virtual worlds and all information representations would be useless if they did not overlap with and share references to our everyday verbal, written and gestural language and activity. "After all, a virtual world filled with virtual offices and virtual desks isn't populated by virtual people, but by real ones."

The differences between media are usually very obvious. The varying physical phenomena and the approaches to connectedness, distinction and (hence) configurationality define each medium as an individuated entity. On the other hand, we tend to overlook their similarities. Their common configurationality makes each medium of information representation potentially semiological, and the hermeneutic circle is a way to describe how symbols in different media can become part of the same language.

Hermeneutic Systems Design

What kind of system design approaches might better support adaptation and appropriation of information representations? Earlier we suggested that, in practice, open implementation is limited in this support because of the degree of control that remains with the programmer. One aspect of information representation that makes it difficult for the programmer to relinquish this control is in the systems of categorisation and abstraction that underlie the system. With computational reflection, this is particularly complex because of the way that information must be considered at many different levels of abstraction, so as to allow the interchangeability of system components at these different levels. We suggested that this ontological complexity adds to the distanciation of the

approach, as deep structure is still relatively 'fixed.' External to the ongoing activity of the users, we see ontological drift.

We propose that there are complementary approaches to system design that reduce these problems. They purposely avoid the fixed categorisations and data abstractions that restrict adaptation, and minimise 'metadata' by relying instead on statistics of usage patterns within a set of symbols. Collaborative filtering and path–based systems involve categorisation that is ephemeral, subjective and, especially with paths, adaptive with the context of use. They aim to minimise abstractions made objectively or *a priori.* They still assume that symbols are objective in that they are contextually independent, shared and persistent. After describing these two approaches, the concluding section will briefly touch upon what such system design approaches might mean for observational CSCW.

Collaborative Filtering

Collaborative filtering (CF) is a burgeoning information access approach based on patterns of subjective ratings of information objects. CF systems show a powerful holism with regard to types of data such as books, films and music. Such 'heterogeneous data' is problematic in traditional database systems and information retrieval (IR) because their metrics of similarity are reliant on uniformly represented object content e.g. the ASCII characters for the word "pipe" are not comparable with a picture of a pipe. Perhaps the earliest published collaborative filtering system was the Tapestry system for accessing eMail and bulletin board messages (Goldberg, 1992). Some of the currently best known research systems are described in (Resnick, 1997), and another widely known system is the book recommender at Amazon.com, a web–based bookstore.

In a basic CF system, recommendations are independent of any current task, query or context. Interaction is driven by building up a profile of ratings. Formalising why we liked or disliked an object is minimised, for example by selecting from a five point rating scale from 'very good' to 'very bad,' and so we can react according to our prejudice. Ratings are thus still categorised, even when simpler schemes of representation, such as a binary scale, are used. Binary scales are often used when more passively obtaining ratings, for example in Amazon.com's book recommender. An action such as purchasing a book is treated as an interpretive act i.e. as giving it a good rating, giving it 1 not 0.

Profiles are matched with each other, based on having similar ratings for the same objects. A set of 'neighbours' (of some fixed and usually small maximum size) is then determined, being the most similar profiles to one's own. This is sometimes described as a way of finding similar people. (The first collaborative filtering system is said to have been an automatic dating agency.) Ordinarily, however, the system presents to you a small amount of information that similar people rated highly but that you have not yet rated. Sometimes one can explicitly choose the subset of people from that profiles should be drawn, allowing use of

knowledge of the user community.

If you rate one or more recommendations, you feed back into the process of adaptation of your profile, your neighbour set, and your neighbours' neighbour sets. Ripples of change spread out, demonstrating adaptivity and configurationality. The hermeneutic circle is strong here, and there is no meaning for a rated object independent of a person's interpretation. If an object has no ratings, it can not be recommended. Initial ratings must be generated manually, through passive rating of actions in another information access system (e.g. book purchases, above) or other means external to the CF system.

Sets of ratings are useful objects in themselves, and system use involves a kind of sharing, but generally they are not put in common for view, discussion and more direct manipulation. Occasionally, however, we can treat a rating and its contents as symbols, holistically. Many people, for example, have a web page of 'hot links' that lists the web pages they particularly like. If this page was instead an active representation of a user's profile in a CF system, and the URL for this page was amongst the symbols manipulated within the system, then we would see metadata slipping down to being data again. Accessing such a page might suggest a passive rating of the pages it references.

Exploratory information access can be driven by the choice of the people involved in recommendation, but exploration is generally constrained to rating objects in ways that will shift and change your profile. What the system treats as relevant for presentation is dependent on all past ratings, and not just the most recent ratings or other current activity. There is no equivalent of the IR query that explicitly states the current information need. Taking account of the temporal order of activity is at the core of an extension of collaborative filtering, the path model.

The Path Model

The path model reflects a desire to take account of a wider range of phenomena and a more contextually driven notion of interaction than CF. A means to do this was made by direct analogy with a theory of urban structure, use and development: the 'space syntax' of (Hillier, 1996). Space syntax puts movement and visibility at the centre of city structure and development. It uses people's paths through the city as expressions of their activities, interests, and associations. Hillier deliberately avoids *a priori* categorisation of urban spaces, instead relying on building up statistics of movement and activity by the public. Analysis is in terms of the configuration of buildings and streets rather than the content or functionality of individual urban elements. He emphasises the importance of considering the extended paths people take rather than only transitions between pairs of city elements. The analogy with information representation relies on treating histories of information use like paths through the city, treating written symbols like spatial forms, and language as the city. This analogy itself has a long

history, with one notable landmark being in the centre of Wittgenstein's language games (1958:8). A more detailed discussion of the analogy with space syntax and the relationship between informatics, architecture and language is given in (Chalmers, 1999b). Although the recency of the path model suggests that some technical detail might be useful in this section, fuller implementation details of prototype systems for recommending and visualising heterogeneous data are described in (Chalmers, 1998) and (Chalmers, 1999a).

A *path* is a time–ordered history of a person's use of symbols. It serves as the representation of user activity. Observed use of a symbol by a person is treated as an interpretive act, and logged. The range of symbols currently includes web URLs and the names of local files accessed through the *xemacs* editor, again expressed as URLs. Each path entry is associated with the system's observation of a user accessing a URL in a web browser, or switching between editor buffers. Each user can turn path logging on and off at will. By default, each path is visible to all those who contribute paths i.e. the set of paths is treated a shared resource. At present, content of web pages and files is not recorded in the path.

Context or current activity is taken into account, in that each instance of a symbol's use is associated with that person's temporally close path entries. We then use patterns of symbol recurrence in determining relevance, by treating the most recent sequence of path entries as an implicit request for recommendations. Every few minutes, the system takes the symbols most recently used i.e. the end of the path, and then searches for past occurrences of each symbol. This search can be among a number of paths within the shared path set i.e. users can select which paths (or path owners) to draw recommendations from, letting them use knowledge of their colleagues to guide the recommendation process. The system then collects the context of each past occurrence of the most recently used symbols—each past occurrence's temporally close path entries. It then collates these sequences, finds which symbols most frequently occur in them, and then presents this ranking as a recommendation list to the user. (An example of a ranked list of symbols that appear to be relevant to the most recently logged activity is shown in Figure 1 and discussed below.) The people whose paths contributed to recommendations are not identified in this list (even though no–one as yet has objected to the idea of being so identified). The system thus recommends to a person symbols that were frequently used in similar contexts but that it has not observed recently in that person's path. 'Similarity' here is interpreted as statistically consistent co–occurrence in observed use of symbols.

The list shown in Figure 1 was made for the author while he was resident in Zürich, Switzerland. The example is intended to demonstrate how recommendations, such as ski information given weather pages, might not be obviously useful until one considers the user's particular context and the history behind it. If the mountain weather was good, the author might ski. One checked the weather before getting details of piste conditions and cable car times. In this

```
http://www.arosabergbahnen.ch/                                  │18.0
http://www.arosabergbahnen.ch/Grafiken/collage.jpg              │15.0
http://www.klosters.ch/images/tn_gotschna.jpg                   │9.0
http://www.klosters.ch/gotschna.html                            │9.0
http://194.158.230.224:9090/telenet/CH/180/2.g-html             │9.0
http://www.arosa.ch/main.html                                   │9.0
http://www.arosa.ch/skiauswahl.html                             │9.0
http://ad.adsmart.net/src/goski/mountains^1?adtype=ac&bgcolor=F...│9.0
http://www.arosabergbahnen.ch/home.html                         │9.0
http://www.arosa.ch/                                            │9.0
```

Figure 1. A recommendation list from an early prototype path system. Starting to choose a day for a ski trip, the author accessed web pages with detailed weather reports for the mountains of Switzerland, including the *telenet* service of a local university. Recommended URLs were drawn from six sequences of past activity in three people's paths, and were mostly for web sites of ski resorts near Zürich, such as Arosa and Klosters. The numbers in the right hand column are weighted tally values used in a rough recommendation ranking.

case, the author got the Klosters information from his own path, and the Arosa information from a colleague's path. Never having been to Arosa, or to the Arosa web site, the recommendation was therefore both novel and relevant. It reflects patterns of activity that are particular to people from that part of Switzerland, and who share interests and activities related to skiing. The example also offers a contrast with traditional database/IR tools which, given as input recently accessed weather pages, would most likely recommend (or retrieve) still more weather pages. Lastly, the recommendations include mixed data types: JPEG images as well as pages of HTML. This demonstrates the ability to handle media usually handled disjointly.

A mouse click on a recommended URL triggers access and display in the web browser. In this one–dimensional list, it is however difficult to gain an impression of how the recommended symbols relate to each other, and so we offer a 2D 'map–like' visualisation of the sequences of symbols collated in the recommendation process. (Details of these are left to a forthcoming paper.) Recommendations are made every few minutes, generally take less than a minute to make, and we can create a matching visualisation in approximately one minute.

While recurrence statistics form categorisations or formalisations, path entries and their patterns are not interpreted or categorised *a priori.* Instead, this is done anew for each person at the time of, and using the context of, each recommendation operation. Each new extension to a path changes the pattern of symbol co–occurrences throughout the shared set of paths. Even if you have not accessed new information recently, your recommendations may still vary as other people's activities change the path configuration. Thus a path system adapts with every use, strengthening the hermeneutic circle. The 'meaning' of a symbol is determined by its pattern of occurrences, and hence its pattern of co–occurrence with other symbols. Symbols are thus represented configurationally. There is no meaning for a symbol independent of paths and, consequently, of individuals' interpretation and use.

We should consider paths as objects within the path model, for example

recording whose paths we select in guiding the recommendation process. If a path has a unique and persistent identifier, such as the name of a file where it is permanently stored, then we can pull this meta–information down to be information within the system.

Complementary System Design Approaches

We see hermeneutics as a useful tool in changing system design practice at a fundamental level—as is the goal of technomethodology. However, we take a somewhat complementary approach to the design of rich, adaptive and responsive representations of information. Where open implementation and computational reflection focuses on complex architectures involving causally connected, formalised levels of abstraction, in our path model work we build on ephemeral correlations. As soon as a model of current activity leads to a symbol being recommended and then used, that model is invalidated and discarded. Unlike reflection's heavy design investment prior to program execution, we use very simple designs that take advantage of the computer being a cheap and tireless observer of activity, albeit one of limited acuity. We cannot avoid the fact that our programs are formal representations, and we still fix in advance how paths are made, manipulated and compared, but we try to minimise the scale and scope of our formalisation—and plan to do more in this regard.

Although it may at first seem like a weakness, another reason the path model offers promise is the minimal functionality a path system provides for each recommended or visualised symbol—usually only simple presentation to the user, or presentation to another tool e.g. passing a URL to a web browser for download and display. Consequent interpretation or activity—on the part of the user or a receiving tool—is not our system's responsibility. It never can completely be, as no program can represent all of the relevant activity. Instead we rely on our system being able to observe a useful subset of the symbols arising from that activity. The path model thus depends on external interpretation and action, which leads to the introduction of new symbols to the configuration.

Ricoeur pointed out that positivist approaches, such as traditional information retrieval, allow access to open semiological systems while structuralist approaches, such as the path model, are dependent on a closed system. While positivism cannot grasp the complex interrelations of symbols, it provides an essential means to access new information for which no interpretation is yet available. Our recommendation system would be useless without people's use of Web search engines, mail tools, magazines, and so forth. Again, we see this not as a weakness, but as realistic acceptance of the holism of tools. We interleave such tools in everyday use, and the sets of objects (or symbols) they operate on overlap. For example, a URL of interest may first be seen in a mail message or a report in a word processor but, at present, if the URL is accessed again then information access systems do not remind us of the mail message or report. Tools

are interrelated by their strengths, weaknesses, similarities and differences in the same way that the data they work on are interrelated. Rather than treat tools in isolation, we see integration as a way to gain the greater power of the ensemble or configuration of parts. We can treat both tools and data semiologically and, meta–data being data, within the same system. Path–based tools should log the names of the tools they are used with, so as to expand the range of recorded contextual phenomena and, for example, to recommend to a user tools that may be unfamiliar, but have been found useful in similar contexts by others.

One might consider that the 'meta–level' of information access tools is the level of discourse on information access approaches and informatics itself. Although we might say that we observe activities at this level, perhaps such observation is not something computers do. However, our holistic view can reach this level. If, when we consider the everyday use of tools on our own computers, we see tools' use as being strongly interwoven, then their supporting areas of discourse are interrelated. In this case, practice appears to be ahead of theory. Web browsers are expanding to become 'integrated desktops,' interconnecting browsers, mail tools, editors, spreadsheets and databases. However, in mainstream information retrieval, for example, there is virtually no theoretical work that explicitly takes account of the fact that IR and database systems are now predominantly used in combination with mail tools and other desktop components, rather than in 70s–style 'batch mode' isolation.

Conclusion: Observations and Metatheory

The path model is a move forward from collaborative filtering, towards a practical approach to observing and recording the socially–held rules used to generate user activity i.e. the utterance. 'Rules' here refers to statistically consistent patterns of symbol use. Compared to traditional informatics, CF and paths offer a contrasting approach to representing and providing information, and are more closely aligned with contemporary philosophy, as exemplified here by hermeneutics. Such systems are, at present, limited in ways they can record and combine events and actions. They are naïve and simple compared to the human observer, but they are still young. Such patterns are not, *a priori*, directed towards any particular type of account or abstraction of activity, as ethnomethodology focuses on agency in the production of society. They evolve with activity, and can run for indefinitely long periods of time—for periods impractically long for human observers. These 'hermeneutic systems' take a shorter route in supporting activity, cutting out the extraction of rules by a sociologist, to convey to a technologist, to build programs finally inserted back into the original setting. While this latter role might be seen as 'servant to the technologist'—rather unlikely to be the main motivation for ethnomethodologists—it is another potential example (positive or negative) of how such systems can do some of the

work that ethnomethodologists do. It is too early to do much more than speculate. Nevertheless, in accord with this paper's emphasis on holism and interdependence, we suggest that an understanding of what such systems can and cannot do, and critically aware use of their results, may mean that they become a useful addition to the sociologist's toolkit of techniques and devices.

Within CSCW, the notion that informatics is just one part of a whole—one of the fields that make up our shared toolkit of techniques and devices—is well accepted. In this paper we have tried to add hermeneutics to this toolkit, or at least to increase our awareness of hermeneutics' utility. We have also highlighted common features of CSCW theory and system design practice, and we see the possibility for hermeneutics to be a unifying meta–theory for CSCW. The intellectual tools of CSCW form a common information space within which we debate and adapt our views on such topics as the phenomena each field handles, the rules of theory and patterns of practice that serve as maps or scripts for our research work, how we can adapt and improve these rules and patterns, what the most relevant parts of the CSCW 'toolset' are to apply when we are studying a particular situation, and how each one interrelates and is interdependent with others. Here, in our point of view, we work on hermeneutics. We adapt and strengthen it, testing its claim to universal applicability, with each turn of our hermeneutic circle.

References

Bannon, L. & Bødker, S. (1997) 'Constructing Common Information Spaces', *Proc. ECSCW 97*, 81–96.

Bentley, R. & Dourish, P. (1995) 'Medium versus Mechanism: Supporting Collaboration Through Customisation', *Proc. ECSCW 95*, 133–148.

Button, G & Dourish, P. (1996) 'Technomethodology: Paradoxes and Possibilities', *Proc. ACM CHI 96*, 19–26.

Chalmers, M., Rodden, K. & Brodbeck, D. (1998) 'The Order of Things: Activity–Centred Information Access', *Proc. WWW7,* 359–367.

Chalmers, M. (1999a) 'Comparing Information Access Approaches,' to appear *in J. American Society for Information Science (50th Volume Special Edition)*.

Chalmers, M. (1999b) 'Informatics, Architecture and Language', to appear in *Social Navigation: Footprints in the Snow*, Munro, A., Hook, K. & Benyon, D. (eds.), Springer, 1999.

Coyne, R. (1995) *Designing Information Technology in the Postmodern Age*, MIT Press.

Dourish, P. (1995) 'Developing a Reflective Model of Collaborative Systems', *ACM Trans. CHI*, 2(1), 40–63.

Dourish, P. (1997) 'Reconsidering Software Representations', *Proc. Workshop on Representations in Interactive System Design*. Also available from the paper's author as: http://www.parc.xerox.com/csl/members/dourish/papers/representation.html

Gadamer, H.–G., (1989) *Truth and Method*, 2nd edn., trans. J. Weinsheimer & D. Marshall, Crossroad.

Giddens, A. (1995) 'Garfinkel, Ethnomethodology and Hermeneutics', in *Politics, Sociology and Social Theory: Encounters with Classical and Contemporary Social Thought*, Stanford University Press, 233–245.

Goldberg, D., et al. (1992) 'Using Collaborative Filtering to Weave an Information Tapestry', *Comm. ACM* 35(12), 61–70.

Grondin, J. (1994) *Introduction to Philosophical Hermeneutics,* trans. J. Weinsheimer, Yale University Press.

Harrison, S., & Dourish, P. (1996) 'Re–Place–ing Space: The Roles of Place and Space in Collaborative Systems', *Proc. CSCW'96.*

Hillier, B. (1996) *Space is the Machine*, Cambridge University Press.

Hurford, J., Studdert–Kennedy, M. & Knight, C. (eds.) (1998) *Approaches to the Evolution of Language*, Cambridge University Press.

Resnick, P. & Varian, H. (1997) (eds.) *Comm. ACM* special issue on Recommender Systems, 40(3).

Ricoeur, P. (1981) *Hermeneutics and Human Sciences: Essays on Language, Action, and Interpretation,* trans. J. Thompson, Cambridge University Press.

Robinson, M. & Bannon, L. (1991) 'Questioning Representations', *Proc. ECSCW 91*, 219–233.

Rorty, R. (1979) *Philosophy and the Mirror of Nature*, Princeton University Press.

de Saussure, F. (1983) *Course in General Linguistics*, trans. Wade Baskin, McGraw–Hill.

Schmidt, K. & Bannon, L. (1992) 'Taking CSCW Seriously: Supporting Articulation Work', *CSCW* 1(1), 7–40.

Schmidt, K. (1997) 'Of Maps and Scripts: The Status of Formal Constructs in Cooperative Work', *Proc. ACM Group 97,* Phoenix, 138–147.

Suchman, L. (1987) *Plans and Situated Actions: The Problem of Human Machine Communication*, Cambridge University Press.

Warnke, G. (1987) *Gadamer: Hermeneutics, Tradition and Reason,* Stanford University Press.

Wegner, P. 'Why Interaction is More Powerful than Algorithms', *Comm. ACM*, 40(5), May 1997, 80–91.

Wittgenstein, L. (1958) *Philosophical Investigations* (3rd ed.), trans. G.E.M. Anscombe, Oxford University Press.